

Computation with narrow CTCs^{*}

A. C. Cem Say and Abuzer Yakaryılmaz

Boğaziçi University, Department of Computer Engineering,
Bebek 34342 İstanbul, Turkey
say,abuzer@boun.edu.tr

January 11, 2013

Abstract. We examine some variants of computation with closed time-like curves (CTCs), where various restrictions are imposed on the memory of the computer, and the information carrying capacity and range of the CTC. We give full characterizations of the classes of languages recognized by polynomial time probabilistic and quantum computers that can send a single classical bit to their own past. Such narrow CTCs are demonstrated to add the power of limited nondeterminism to deterministic computers, and lead to exponential speedup in constant-space probabilistic and quantum computation. We show that, given a time machine with constant negative delay, one can implement CTC-based computations without the need to know about the runtime beforehand.

1 Introduction

It is known [AW09] that adding the capability of sending a polynomial number of bits through a closed timelike curve (CTC) to the past, so that it can be used as part of the input, to models as weak as constant-depth, polynomial-size Boolean circuits increases their computational power significantly, to match that of polynomial-space Turing machines. Interestingly, adding the same capability to a polynomial-time quantum computer results once again in the ability to solve precisely the problems in PSPACE, leading to Aaronson and Watrous’ conclusion [AW09] that “CTCs make quantum and classical computing equivalent.”

Since the information carrying capacity (*width*) of any CTC is finite, and the cost of building such a channel to the past may depend on its width critically, it is important to examine the power of computational models with “narrow” CTCs, i.e. those which are restricted to use a single-bit channel, regardless of the size of the input. The study of narrow CTCs [Bac04,AW09] has focused on polynomial-time computers as the core model until now, with results showing that classical

^{*} This work was partially supported by the Scientific and Technological Research Council of Turkey (TÜBİTAK) with grant 108E142.

computers augmented with narrow CTCs can recognize any language in BPP_{path} , whereas quantum computers with this capability can solve all problems in PP . Since it is not known whether any of the containments in $\text{P} \subseteq \text{NP} \subseteq \text{BPP}_{\text{path}} \subseteq \text{PP} \subseteq \text{PSPACE}$ is proper, we cannot presently say whether narrow CTCs are “useful” in these setups at all, and if so, whether they confer the same amount of power to classical and quantum models.

In this paper, we will complement the results mentioned above to provide full characterizations of the classes of languages recognized by probabilistic and quantum computers with narrow CTCs. It turns out that a CTC with a single classical bit provides a computer with precisely the power of postselection [Aar05].

The effects of restrictions on the core model of computation to be used inside the CTC on the power of the resulting setup were posed as open questions in [AW09]. In this regard, we consider real-time probabilistic and quantum finite automata (PFAs and QFAs), as well as several deterministic automaton models, as possible core models to be augmented with a narrow CTC. We show that narrow CTCs add the powers of limited nondeterminism or a certain kind of two-wayness to real-time automata. We are therefore able to prove that real-time PFAs, QFAs and deterministic pushdown automata (DPDAs) with narrow CTCs are strictly more powerful than their standard counterparts, whereas this addition does not increase the power of one-way deterministic finite automata (DFAs), or any deterministic two-way model. We also show that QFAs with narrow CTCs outperform their probabilistic counterparts, in contrast to the Aaronson-Watrous result on polynomial-time computation.

The rest of this paper is structured as follows. Section 2 summarizes previous work on the use of CTCs in computation. We define an alternative model of augmenting computers with CTCs in Section 3. Section 4 establishes the equivalence of the powers of postselection and narrow CTCs. In 5, we examine the effects of endowing weak deterministic machines with narrow CTCs. In Section 6, we demonstrate how our model of Section 3 can be implemented with time machines with constant negative delay. Section 7 is a conclusion.

2 Preliminaries

The existence of CTCs does not seem to be incompatible with the best available theories of spacetime. To cite just one example, an influential paper [MTY88] by Morris, Thorne, and Yurtsever describes how a techno-

logically advanced civilization could first create a “wormhole”, and then transform it to a “time machine” that can be used to send messages, or even people, backwards in time. The time machines of [MTY88] cannot be used to send anything to a time before their date of construction, say, d , and have a constant “range” of, say, T seconds, (determined by their builder at the beginning,) such that a message sent at any time point t ($t > d + T$) is received at time point $t - T$. Once built, the machine can be used as many times as one wishes for such transmissions.

As noted by many authors of science fiction, a major problem with time travel is the “Grandfather Paradox,” where a time traveler from the future prevents himself from traveling in the first place, leading to confusion about the state of the universe at the presumed time, say, A , of his arrival: He arrives if and only if he does not arrive. It was thought that Nature would prevent this logical inconsistency by simply not allowing time travel scenarios of that kind to be realized. Note that this argument assumes that the universe is supposed to be in exactly one, deterministic state, at all times. Probabilistic and quantum theories do not include this restriction, and David Deutsch [Deu91] showed that time travels to the past, including the above-mentioned scenario, would not lead to such problems if one just assumes that Nature imposes a *causal consistency condition* that the state x of the universe in the critical moment should be a fixed point of the operator f describing the evolution in the CTC, i.e. that $x = f(x)$. In the Grandfather Paradox scenario, Nature would “set” the state of the universe at time A to a distribution where the traveler arrives with probability $\frac{1}{2}$ to keep things consistent, as a “response” to the self-prevention action of the traveler.

As Deutsch noted, a computer which sends part of its output back in time to be used as part of its input can solve many computational problems much faster than what is believed to be possible without such loops. Bacon [Bac04] showed that NP-complete problems can be solved by polynomial-time computers with narrow CTCs. Aaronson and Watrous [AW09] proved, as mentioned in the introduction, that $\text{AC}_{\text{CTC}}^0 = \text{P}_{\text{CTC}} = \text{BQP}_{\text{CTC}} = \text{PSPACE}_{\text{CTC}} = \text{PSPACE}$, where the subscript CTC under a class name indicates that the related machines have been reinforced by polynomial-width CTCs.

Let us review Aaronson and Watrous’ model of quantum computation¹ with CTCs from [AW09]: A deterministic polynomial-time algorithm \mathcal{A} takes an input w , and prints the description of a quantum circuit Q_w with rational amplitudes. Q_w acts on two registers of polynomial

¹ The variant where the core model is classical is defined similarly.

size. One register holds the information that is sent from the end of the computation in the future through the CTC, whereas the other one is a standard causality-respecting register, including a bit that will be used to report the output of the computation. The circuit is executed, with the CTC register set by Nature to some state satisfying the causal consistency condition described above, and the causality-respecting register initialized to all zeros, and the result is read off the output bit. A language L is said to be decided by such a CTC algorithm \mathcal{A} if all members of L are accepted with high probability, and all nonmembers are rejected with high probability.

Note that this setup necessitates \mathcal{A} to build a new CTC of the appropriate width for each different input w . This forces one [Deu91] to take the cost of this construction into account when analyzing the complexity, and the resources required may well scale exponentially in the width of the CTC.² The study of narrow CTCs, where this cost does not depend on the input length, is thus motivated.

3 Our model

We will be considering several computation models that are augmented with the capability of sending a single classical bit of information from the time of the end of their execution back to the beginning.

We define a machine of type M_{CTC_1} as simply a machine of type M ,³ which has access to an additional bit in a so-called *CTC cell*. The CTC cell obtains its “initial” distribution from the future, according to the causal consistency condition. The program format of an M_{CTC_1} differs from that of an M so that it specifies the transitions to be performed for all possible combinations of not only the input symbol, internal state, etc., but also the CTC cell value. The set S of internal states of an M_{CTC_1} is defined as the union of three disjoint sets S_n , S_{p_0} , and S_{p_1} . The states in S_n are of the standard variety. When they are entered, the states in S_{p_1} (S_{p_0}) cause a 1 (0) to be sent back in time to be assigned to the CTC cell at the start of the execution. We assume that states in $S_{p_1} \cup S_{p_0}$ are entered only at the end of execution, (for real-time models, this is precisely when the machine is reading the end-marker symbol), and all

² David Deutsch, personal communication.

³ We will examine DFAs, PFAs, QFAs, both real-time and two-way DPDAs, time-bounded probabilistic and quantum Turing machines (PTMs and QTMs), and space-bounded deterministic Turing machines as core models. See [GHI67,Sip06,YS11b] for the standard definitions of these models.

states entered at that time are in $S_{p_1} \cup S_{p_0}$. Any number of members of S can be designated as accept states. The input string w is accepted if, for all stationary distributions of the evolution of the CTC bit induced by w , the machine accepts with probability at least $\frac{2}{3}$ with the CTC cell starting at that distribution. A string w is rejected if, for all stationary distributions of the evolution of the CTC bit induced by w , the machine rejects with probability at least $\frac{2}{3}$ with the CTC cell starting at that distribution. A language is recognized if all its members are accepted and all its nonmembers are rejected.

It is evident that any language recognized by any M_{CTC_1} according to our definition is also decided by some CTC algorithm *à la* the Aaronson-Watrous definition, described in Section 2. The motivation for the difference between the definitions is that the weakness of some of the core models we will use precludes us from performing any processing before using the CTC, and calculating or bounding the runtime, which determines the required “range” of the CTC beforehand. For more on this issue, see Section 6.

4 Postselection and narrow CTCs

An important tool in the analysis of the capabilities of computers with narrow CTCs is the observation that one CTC bit endows any probabilistic or quantum core model with the power of postselection [Aar05]. This fact is already known, but since we have not seen it stated explicitly anywhere, we present a demonstration of it below.

Postselection is the capability of discarding all branches of a computation in which a specific event does not occur, and focusing on the surviving branches for the final decision about the membership of the input string in the recognized language. A formal definition of polynomial time computation with postselection can be found in [Aar05], where it was proven that **PostBQP**, the class of languages recognized by polynomial-time QTMs with postselection, equals the class **PP**. In the following, we present the analogous definition in the context of real-time computation with constant space.

A PFA (resp. QFA) \mathcal{A} with postselection [YS11a] is simply an ordinary PFA (QFA) whose state set is partitioned into the sets of postselection accept, postselection reject, and nonpostselection states, and satisfies the condition that the probability that \mathcal{A} will be in at least one postselection (accept or reject) state at the end of the processing is nonzero for all possible input strings. The overall acceptance and rejection probabilities

of any input string w ($P_{\mathcal{A}}^a(w)$ and $P_{\mathcal{A}}^r(w)$, respectively) are calculated by simply discarding the computational paths ending at nonpostselection states, and performing a normalization so that the decision about the input is given by the postselection states:

$$P_{\mathcal{A}}^a(w) = \frac{p_{\mathcal{A}}^a(w)}{p_{\mathcal{A}}^a(w) + p_{\mathcal{A}}^r(w)} \quad \text{and} \quad P_{\mathcal{A}}^r(w) = \frac{p_{\mathcal{A}}^r(w)}{p_{\mathcal{A}}^a(w) + p_{\mathcal{A}}^r(w)}, \quad (1)$$

where $p_{\mathcal{A}}^a(w)$ and $p_{\mathcal{A}}^r(w)$, respectively, are the acceptance and rejection probabilities of w before the normalization.

Lemma 1. *Any language that can be recognized by a real-time automaton of type $M \in \{PFA, QFA\}$ with postselection can be recognized by an M_{CTC_1} .*

Proof. Let \mathcal{A} be the given automaton with postselection. We construct an M_{CTC_1} \mathcal{A}' . The values 1 and 0 of the CTC bit are associated with acceptance and rejection, as explained below. \mathcal{A}' imitates the behavior of \mathcal{A} until the end of the input. If \mathcal{A} ends up at a postselection state, \mathcal{A}' halts with the same decision as \mathcal{A} at that state. If \mathcal{A} ends up at a nonpostselection state, \mathcal{A}' simply reports the value it sees in the CTC bit as its decision. \mathcal{A}' sends the value associated with its decision to the past as it halts.

The evolution of the CTC bit of \mathcal{A}' for input w is described by the column stochastic matrix

$$\begin{pmatrix} 1 - p_{\mathcal{A}}^a(w) & p_{\mathcal{A}}^r(w) \\ p_{\mathcal{A}}^a(w) & 1 - p_{\mathcal{A}}^r(w) \end{pmatrix}, \quad (2)$$

whose only stationary distribution is

$$\begin{pmatrix} \frac{p_{\mathcal{A}}^r(w)}{p_{\mathcal{A}}^a(w) + p_{\mathcal{A}}^r(w)} \\ \frac{p_{\mathcal{A}}^a(w)}{p_{\mathcal{A}}^a(w) + p_{\mathcal{A}}^r(w)} \end{pmatrix}, \quad (3)$$

where the first and second entries stand for the probabilities of the values 0 and 1, respectively, meaning that \mathcal{A}' recognizes \mathcal{A} 's language with the same error probability. \square

Let $\text{BPP}_{\text{CTC}_1}$ and $\text{BQP}_{\text{CTC}_1}$ denote the classes of languages recognized by polynomial time PTMs and QTMs with narrow CTCs using classical bits, respectively.⁴ BPP_{path} is the class of languages recognized by polynomial-time PTMs with postselection.

⁴ Our definition of $\text{BQP}_{\text{CTC}_1}$ is different from the $\text{BQP}_{\text{CTC}_1}$ given in [AW09], since Aaronson and Watrous consider a quantum bit sent through the CTC.

The results $\text{BPP}_{\text{path}} \subseteq \text{BPP}_{\text{CTC}_1}$ and $\text{PP} \subseteq \text{BQP}_{\text{CTC}_1}$, that we alluded to in the introduction, are obtained using the link described above between CTCs and postselection.

We now present our main result that the power of postselection is all that a narrow CTC can confer on a computer. Let pPTM and pQTM denote polynomial-time PTMs and QTMs, respectively.

Lemma 2. *Any language that can be recognized by an M_{CTC_1} , where $\text{M} \in \{\text{PFA}, \text{QFA}, \text{pPTM}, \text{pQTM}\}$, can be recognized by a machine of type M with postselection.*

Proof. Let L be a language recognized by an M_{CTC_1} named \mathcal{A} . We start by constructing two machines of the standard type M , namely, \mathcal{A}_0 and \mathcal{A}_1 , that simulate the computation of \mathcal{A} by fixing 0 and 1 for the value of the CTC bit, respectively.

Let $p_{\mathcal{A}_i}^j(w)$ denote the probability that \mathcal{A}_i will reach a configuration corresponding to sending the bit j to the past at the end of its computation when started on input w , where $i, j \in \{0, 1\}$. The CTC bit's evolution is described by

$$\begin{pmatrix} 1 - p_{\mathcal{A}_0}^1(w) & p_{\mathcal{A}_1}^0(w) \\ p_{\mathcal{A}_0}^1(w) & 1 - p_{\mathcal{A}_1}^0(w) \end{pmatrix}, \quad (4)$$

with stationary distribution

$$\begin{pmatrix} \frac{p_{\mathcal{A}_1}^0(w)}{p_{\mathcal{A}_0}^1(w) + p_{\mathcal{A}_1}^0(w)} \\ \frac{p_{\mathcal{A}_0}^1(w)}{p_{\mathcal{A}_0}^1(w) + p_{\mathcal{A}_1}^0(w)} \end{pmatrix}. \quad (5)$$

For $w \in L$, we therefore have the inequality

$$\frac{p_{\mathcal{A}_1}^0(w)}{p_{\mathcal{A}_0}^1(w) + p_{\mathcal{A}_1}^0(w)} p_{\mathcal{A}_0}^a(w) + \frac{p_{\mathcal{A}_0}^1(w)}{p_{\mathcal{A}_0}^1(w) + p_{\mathcal{A}_1}^0(w)} p_{\mathcal{A}_1}^a(w) \geq \frac{2}{3}. \quad (6)$$

We claim that, for all machine types M mentioned in the theorem statement, one can construct an instance of M , say, \mathcal{A}' , which will have two mutually exclusive collections of states, say, S_a and S_r , such that the probability that \mathcal{A}' halts in S_a when started on input w is

$$\frac{1}{2} (p_{\mathcal{A}_1}^0(w) p_{\mathcal{A}_0}^a(w) + p_{\mathcal{A}_0}^1(w) p_{\mathcal{A}_1}^a(w)), \quad (7)$$

and the probability that \mathcal{A}' halts in S_r is

$$\frac{1}{2} (p_{\mathcal{A}_1}^0(w) (1 - p_{\mathcal{A}_0}^a(w)) + p_{\mathcal{A}_0}^1(w) (1 - p_{\mathcal{A}_1}^a(w))). \quad (8)$$

For instance, if $M = \text{pPTM}$, we first build two pPTMs, say, \mathcal{A}_{10} and \mathcal{A}_{01} , for handling the two operands of the addition in Equation 7 by sequencing \mathcal{A}_0 and \mathcal{A}_1 to run on the input in the two different possible orders. \mathcal{A}' simply runs \mathcal{A}_{10} and \mathcal{A}_{01} , with probability $\frac{1}{2}$ each. Equation 8 is handled similarly.

If \mathcal{A} is a real-time machine, the sequential processing described above is not permitted, and we instead perform tensor products of \mathcal{A}_0 and \mathcal{A}_1 to obtain the submachines \mathcal{A}_{10} and \mathcal{A}_{01} .

Once \mathcal{A}' is completed, we view it as a machine with postselection, by postselecting on the states in $S_a \cup S_r$ being reached at the end. S_a will be designated to be the set of accept states. S_r will be the reject states.

Performing the normalization described in Equation 1 on Equations 7 and 8 to calculate \mathcal{A}' 's probability of acceptance, one obtains precisely the expression depicted in Equation 6. The case of $w \notin L$ is symmetric. We conclude that \mathcal{A}' recognizes L with exactly the same error probability as \mathcal{A} . \square

We have therefore proven

Theorem 1. *For any $M \in \{PFA, QFA, pPTM, pQTM\}$, M_{CTC_1} is equivalent in language recognition power to a machine of type M with postselection.*

Corollary 1. $\text{BPP}_{\text{path}} = \text{BPP}_{\text{CTC}_1}$, and $\text{PP} = \text{PostBQP} = \text{BQP}_{\text{CTC}_1}$.

We can use Lemma 1 to demonstrate the superiority of real-time PFAs and QFAs with narrow CTCs over their standard versions, which can only recognize regular languages with bounded error: For a given string w , let $|w|_\sigma$ denote the number of occurrences of symbol σ in w . The nonregular language $L_{eq} = \{w \in \{a, b\}^* \mid |w|_a = |w|_b\}$ can be recognized by a PFA with postselection [YS11a]. As for quantum machines, the language $L_{pal} = \{w \in \{a, b\}^* \mid w = w^r\}$ is recognized by a QFA with postselection [YS11a]. L_{eq} is known [Fre81] to be recognizable by two-way PFAs at best in superpolynomial time [GW86], and the best known two-way QFA algorithm [AW02] for L_{pal} has exponential expected runtime. Furthermore, L_{pal} is known [DS92] to be unrecognizable by even two-way PFAs with bounded error, and no PFA with postselection can outperform a standard two-way PFA [YS11a], so we have established that finite-state quantum models with narrow CTCs outperform their probabilistic counterparts:

Corollary 2. *The class of languages recognized by QFA_{CTC_1} s properly contains the class of languages recognized by PFA_{CTC_1} s.*

5 Weak deterministic models with narrow CTCs

We adapt the argument used in [AW09] to prove that $P_{CTC} \subseteq PSPACE$ to state the following upper bounds for the powers of deterministic machines with narrow CTCs:

Theorem 2. *Let \mathcal{A} be any machine of type M_{CTC_1} that recognizes a language according to the definition in Section 3, where M is a deterministic model. \mathcal{A} can be simulated by running two machines of type M in succession.*

Proof. Assume that \mathcal{A} never enters an infinite loop. (We get rid of this assumption in Section 6.) For a given input string w , let f_w be the mapping among probability distributions over the CTC bit realized by \mathcal{A} when running on w . We have to find a distribution d , such that $f_w(d) = (d)$, and see how \mathcal{A} responds to w when it starts with the CTC bit having that distribution.

We first run a machine \mathcal{A}_1 of type M obtained by fixing the CTC bit of \mathcal{A} to 0 on input w . If this run ends at a state belonging to S_{p_0} , the collection of \mathcal{A} 's states that set the CTC bit to 0, we have found that $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ is a stationary distribution, and the response given to w by \mathcal{A}_1 is what \mathcal{A} itself would announce if it were executed. If this first stage ends within S_{p_1} , then we run another machine \mathcal{A}_2 of type M , obtained by fixing the CTC bit of \mathcal{A} to 1, on input w . Note that the only remaining possibilities for stationary distributions at this stage are $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix}$, and in either case, \mathcal{A}_2 's response to w is certain to be identical to \mathcal{A} 's response, since \mathcal{A} cannot have an error probability as big as $\frac{1}{2}$. \square

This construction can be realized by a two-way version of model M . It is known [She59] that two way DFA's are equivalent to their one-way versions.

Corollary 3. *One-way DFA_{CTC_1} s recognize precisely the regular languages.*

Two-way DPDAs are more powerful than one-way DPDAs [GHI67]. Given a one-way $DPDA_{CTC_1}$ \mathcal{A} , we can apply the idea of Theorem 2 to obtain three DPDAs as follows: \mathcal{A}_1 is obtained by fixing the CTC bit to 0, and accepting if computation ends in a member of S_{p_0} that is also an accept state. \mathcal{A}_2 is obtained by fixing the CTC bit to 0, and accepting

if computation ends in any member of S_{p_1} . \mathcal{A}_3 is obtained by fixing the CTC bit to 1, with no change to the accept states of \mathcal{A} .

Calling the languages recognized by these three machines L_1 , L_2 , and L_3 , respectively, it is easy to see that the language recognized by \mathcal{A} is $L_1 \cup (L_2 \cap L_3)$. We conclude that any language recognized by a one-way $DPDA_{CTC_1}$ can be expressed as the union of a deterministic context-free language (DCFL) with a language that is the intersection of two DCFLs.

To demonstrate that $DPDA_{CTC_1}$ s are actually more powerful than ordinary DPDAs, we will show that the capability of sending a finite number of bits to the past endows a machine with the power of limited nondeterminism.

The amount of nondeterminism used by a PDA can be quantified in the following manner [Her97]: The branching of a single move of a (non-deterministic) PDA is defined as the number of next configurations that are possible from the given configuration. The branching of a computation path of a PDA \mathcal{N} is the product of the branchings of all the moves in this path. The branching of a string w accepted by \mathcal{N} is the minimum of the branchings of the paths of \mathcal{N} that accept w . Finally, the branching of \mathcal{N} is the maximum of the branchings of the strings accepted by \mathcal{N} .

Fact 1. *The class of languages recognized by PDAs with branching k is the class of unions of k DCFLs.*

Theorem 3. *Any language that can be expressed as the union of two DCFLs can be recognized by a one-way $DPDA_{CTC_1}$.*

Proof. By Fact 1, we only need to show how to build a $DPDA_{CTC_1}$ that can simulate a given PDA \mathcal{N} with branching 2.

We convert \mathcal{N} to an equivalent PDA \mathcal{N}' , all of whose computational paths have branching exactly 2, by modifying the program so that for any computational path of \mathcal{N} with branching greater than 2, \mathcal{N}' simply scans the input until the end, and rejects without performing that excess branching. For every nondeterministic state of \mathcal{N}' , name the two outgoing branches 0 and 1. Convert \mathcal{N}' to a $DPDA_{CTC_1}$ \mathcal{A} which simulates \mathcal{N}' , selecting the i th branch if and only if it sees the value i in the CTC bit. At the end of the input, \mathcal{A} sends the name of the current branch to the past if it is accepting the input. It sends the name of the other branch otherwise.

We consider the four possible cases corresponding to accept/reject responses of the two paths, and the associated evolutions of the CTC bit in Table 1. It is evident that \mathcal{A} recognizes the language of \mathcal{N} with zero error. \square

Table 1. Evolutions and stationary distributions of the CTC bit of \mathcal{A}

branch ₀	branch ₁	CTC transformation	Stationary distribution
Acc	Acc	$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$	Any distribution
Acc	Rej	$\begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 \\ 0 \end{pmatrix}$
Rej	Acc	$\begin{pmatrix} 0 & 0 \\ 1 & 1 \end{pmatrix}$	$\begin{pmatrix} 0 \\ 1 \end{pmatrix}$
Rej	Rej	$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$	$\begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix}$

Since there exist languages (e.g. $\{a^i b^j c^k | i = j \text{ or } i = k\}$) that are not deterministic context-free, but which can be expressed as the union of two DCFLs, we conclude that the computation power of DPDAs is actually increased by the addition of a narrow CTC.

6 Implementation of arbitrary negative delays using a constant-length channel

In our model of computation with CTCs, the runtime, (that is, the length of the backward jump in time that the CTC bit will go through,) is not known at the start of the computation. In the Aaronson-Watrous model, this “range” is known at the end of the preprocessing stage prior to the point where the information arrives through the CTC, so one can view their program as building a mechanism that can transport a specified quantity of information with that fixed range, and then using it. Let us show that our setup can be implemented using a time machine of the Morris-Thorne-Yurtsever variety⁵ of constant range T seconds.

Assume that all instructions in our programming language take an equal amount of time to be executed, and that a backward jump of T seconds amounts to a difference of $k > 3$ instructions, in the sense that if the i th instruction sends a bit backwards, that bit will be available to the $(i-k)$ th, but not to the $(i-k-1)$ th instruction.

Consider the trace of a computation conforming to our definition of Section 3 in Figure 1. In the figure, the first instruction r indicates the point where the CTC bit coming from the future is received, and s indicates the point where it is sent. (For simplicity, assume that the program idles (performs a no-op) whenever an r “instruction” is executed in this and the following figure.) The idea in transforming this program to one which uses a fixed-length channel to the past is simply for the information

⁵ See the brief description in Section 2.

to travel from the end of the execution back to its start in several “hops,” rather than one. For this purpose, we rewrite our programs so that they send back the current value of the CTC bit backwards once every $k - 1$ instructions. The modified program would then have a computation of the kind depicted in Figure 2.



Fig. 1. Trace of a computation with a variable-length channel

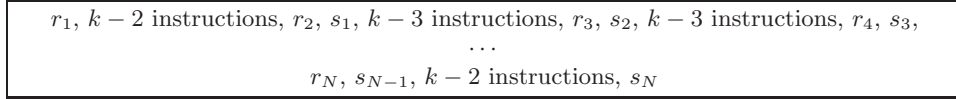


Fig. 2. Trace of a computation with a fixed-length channel

In Figure 2, the r_i - s_i pairs indicate the corresponding arrival and departure points of the CTC bit. The modification to be performed on the program is clearly achievable by a finite-state mechanism. Note that the first and last “segments” (i.e. sequences of standard instructions, not involving time travel) of the execution are one step longer than the intermediate segments, since we do not need send or receive to or from time points outside the duration of the execution. Also note that the last segment has to be padded with no-ops if necessary⁶ to make sure that r_N and s_N are separated with the proper length of time.

We show that the fixed- vs. variable-length channel setups are equivalent, by demonstrating that the application of the causal consistency condition to all r_i - s_i pairs of Figure 2 yields the same restriction on stationary distributions as the single application in Figure 1.

We will assume, without loss of generality, that our original M_{CTC_1} program consults its CTC bit only at the beginning of its computation. Note that this program can now be considered to consist of two branches

⁶ This necessity, and the inclusion of the intermediate r and s instructions that do not consume input symbols, mean that the resulting program is not real-time, although we can insert paddings in our language definitions of earlier sections to obtain sister languages that can be recognized in real-time by the approach of this section.

b_0 and b_1 , differentiated by the initial CTC value. Let p^i denote the probability that a 0 will be sent back at the end of the execution, given that branch b_i has been selected at the beginning. It is important to see that these values are unchanged by a switch between the two setups.

Consider all pairs of the form $(r_{i-s_i}, r_{i+1-s_{i+1}})$, where $1 < i < N$ in one of these two branches. Let the stationary distribution associated with the r_{i-s_i} pair be $\begin{pmatrix} p_i \\ 1 - p_i \end{pmatrix}$. This must equal the stationary distribution of the $r_{i+1-s_{i+1}}$ pair, since the value sent back by s_i equals the one received by r_{i+1} : The evolution in the r_{i-s_i} pair is

$$\begin{pmatrix} p_{i+1} & p_{i+1} \\ 1 - p_{i+1} & 1 - p_{i+1} \end{pmatrix}, \quad (9)$$

whose only stationary distribution is indeed $\begin{pmatrix} p_{i+1} \\ 1 - p_{i+1} \end{pmatrix}$.

Now consider the evolution in the r_1-s_1 pair. Using superscripts again to denote the probabilities associated with the two branches, this is

$$\begin{pmatrix} p_2^0 & p_2^1 \\ 1 - p_2^0 & 1 - p_2^1 \end{pmatrix}, \quad (10)$$

which, by the fact that $p_2^0 = p_N^0$ and $p_2^1 = p_N^1$, equals

$$\begin{pmatrix} p_N^0 & p_N^1 \\ 1 - p_N^0 & 1 - p_N^1 \end{pmatrix}. \quad (11)$$

Since p_N^0 and p_N^1 are exactly the probabilities of 0 being assigned to the CTC bit by the respective branches in the variable-length channel setup as well, the matrix above is precisely the evolution matrix induced by that program running on the same input.

One useful consequence of this result is that it enables us to consider core models in which the machine can have multiple computation paths of different length in a single run. (Note that the models of Sections 4 and 5 avoided this issue, since they were either restricted to real-time processing, or had sufficient resources to ensure that all paths have the same runtime, with any premature acceptances or rejections postponed until the scanning of the right input-end-marker.) In the Aaronson-Watrous model, this possibility can be handled by effectively making sure, by adding no-ops to short paths, that all paths have the same runtime, and building the CTC according to that specification. For the weak models we are considering, it is sufficient to view the CTC-assisted computation as being

implemented with fixed-length channels. Note that the two branches of the computation need not be of equal length for the reasoning above to hold. This will help us characterize the power of certain models with two-way access to the input, as soon as we clarify the issue of the possibility of some branches entering infinite loops.

What happens if a M_{CTC_1} program (for a deterministic \mathcal{M}) employing a fixed-length channel enters an infinite loop on, say, branch 0, never reaching accept or reject states? It is of course more realistic to assume that in this case the execution will finish, (possibly a very long time later), due to external reasons. In any case, no new value will be assigned to the CTC bit on that branch. The reasoning above about all the p_i^0 equaling each other for $i > 1$ still applies, but since no assignment is performed within the last hop, consideration of the end of the branch does not help us to constrain this probability value. Assuming that branch 1 halts, sending 0 to the past with probability p^1 , the evolution within the r_1 - s_1 period is

$$\begin{pmatrix} 1 & p^1 \\ 0 & 1 - p^1 \end{pmatrix}. \quad (12)$$

$\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ is a stationary distribution of this matrix. The program neither accepts nor rejects when the CTC bit starts with this distribution, and therefore such programs, i.e. those with infinite loops, do not contribute to the class of languages recognized by M_{CTC_1} s. Since M_{CTC_1} programs with finite-length branches can be handled using fixed-length channels, and simulated using the technique of Theorem 2, we have

Corollary 4. *For any space bound s , $SPACE(s)_{CTC_1} = SPACE(s)$.*

Corollary 5. *The computational power of two-way DPDAs does not change with the addition of a narrow CTC.*

7 Concluding remarks and open questions

We have examined the power of several computational models augmented by the capability of sending a single classical bit to the past. We have characterized the classes BPP_{CTC_1} and BQP_{CTC_1} in terms of classical conventional classes, and shown that real-time probabilistic and quantum finite automata, as well as deterministic pushdown automata, benefit from narrow CTCs. In Section 6, we establish that CTCs remain useful even if the information channel to the past has a small fixed range, e.g. a few

seconds, and that narrow CTCs do not change the power of deterministic models with two-way access to the input string. One-way deterministic finite automata are also not affected.

In an earlier paper [YS11a], we had shown that machines with postselection have precisely the same power as conventional machines that are able to reset their input head to the start of the input string and switch to the initial state, to restart the computation all over again [YS10]. The new link to narrow CTCs shows that postselection is indeed a profoundly interesting concept that requires further investigation.

Some open questions remain. Can we pin down the power of $DPDA_{CTC_1}$ s further than we have done in Theorems 2 and 3? Most of our results are obviously generalizable to CTCs with capacities of $k > 1$ bits. With more CTC bits, one can clearly implement more nondeterministic choices. With more nondeterminism, one can obtain more succinct finite automata for certain languages, and build PDAs with superior language recognition capability [Her97]. It would be interesting to clarify and quantify these relationships.

In most of our programs, for instance, the machines of Lemma 1, the machine’s decision about the input can be read off the CTC bit immediately after the start of the execution with high probability, without the need to wait for the whole input string to be scanned. What stops us from “cheating,” turning the computer off and using the output “for free,” at that point? The answer, according to the analysis in [Deu91], is that the CTC bit will be set to the desired distribution only if nothing interferes with the complete scheduled run of the machine, and its value will be “corrupted” if, say, a meteorite is likely to hit the computer in the middle of the computation, or if the molecules in our brains are configured so that we are likely to cheat. This requirement of perfect isolation from external interference during the computation is yet another interesting aspect of CTC-assisted computation.

Acknowledgements

We thank David Deutsch, Scott Aaronson, Amos Ori, and Taylan Cemgil for their helpful answers to our questions.

References

- Aar05. Scott Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proceedings of the Royal Society A*, 461(2063):3473–3482, 2005.

- AW02. Andris Ambainis and John Watrous. Two-way finite automata with quantum and classical states. *Theoretical Computer Science*, 287(1):299–311, 2002.
- AW09. Scott Aaronson and John Watrous. Closed timelike curves make quantum and classical computing equivalent. *Proceedings of the Royal Society A*, 465(2102):631–647, 2009.
- Bac04. Dave Bacon. Quantum computational complexity in the presence of closed timelike curves. *Physical Review A*, 70(032309), 2004.
- Deu91. David Deutsch. Quantum mechanics near closed timelike lines. *Physical Review D*, 44(10):3197–3217, November 1991.
- DS92. Cynthia Dwork and Larry Stockmeyer. Finite state verifiers I: The power of interaction. *Journal of the ACM*, 39(4):800–828, 1992.
- Fre81. Rūsiņš Freivalds. Probabilistic two-way machines. In *Proceedings of the International Symposium on Mathematical Foundations of Computer Science*, pages 33–45, 1981.
- GHI67. James N. Gray, Michael A. Harrison, and Oscar H. Ibarra. Two-way push-down automata. *Information and Control*, 11:30–70, 1967.
- GW86. Albert G. Greenberg and Alan Weiss. A lower bound for probabilistic algorithms for finite state machines. *Journal of Computer and System Sciences*, 33(1):88–105, 1986.
- Her97. Christian Herzog. Pushdown automata with bounded nondeterminism and bounded ambiguity. *Theoretical Computer Science*, 181:141–157, 1997.
- MTY88. Michael S. Morris, Kip S. Thorne, and Ulvi Yurtsever. Wormholes, time machines, and the weak energy condition. *Physical Review Letters*, 61(13):1446–1449, September 1988.
- She59. John C. Shepherdson. The reduction of two-way automata to one-way automata. *IBM Journal of Research and Development*, 3:198–200, 1959.
- Sip06. Michael Sipser. *Introduction to the Theory of Computation*, 2nd edition. Thomson Course Technology, United States of America, 2006.
- YS10. Abuzer Yakaryılmaz and A. C. Cem Say. Succinctness of two-way probabilistic and quantum finite automata. *Discrete Mathematics and Theoretical Computer Science*, 12(4):19–40, 2010.
- YS11a. Abuzer Yakaryılmaz and A. C. Cem Say. Probabilistic and quantum finite automata with postselection. Technical Report arXiv:1102.0666, 2011.
- YS11b. Abuzer Yakaryılmaz and A. C. Cem Say. Unbounded-error quantum computation with small space bounds. *Information and Computation*, 279(6):873–892, 2011.